TRANSMITTAL OF APPEAL BRIEF (Large Entity)	Docket No. 13845

In Re Application Of: **Peter N. Yianilos et al.**

Application No.	Filing Date	Examiner	Customer No.	Group Art Unit	Confirmation No.
09/927,589	August 10, 2001	Etienne Pierre Leroux	23389	2161	8179

Invention: **SYNCHRONIZABLE TRANSACTIONAL DATABASE METHOD AND SYSTEM**

COMMISSIONER FOR PATENTS:

Transmitted herewith is the Appeal Brief in this application, with respect to the Notice of Appeal filed on **September 29, 2005**

The fee for filing this Appeal Brief is: **$500.00**

☒ A check in the amount of the fee is enclosed.

☐ The Director has already been authorized to charge fees in this application to a Deposit Account.

☒ The Director is hereby authorized to charge any fees which may be required, or credit any overpayment to Deposit Account No. **19-1013/SSMP**

☐ Payment by credit card. Form PTO-2038 is attached.

WARNING: Information on this form may become public. Credit card information should not be included on this form. Provide credit card information and authorization on PTO-2038.

Signature

Dated: **December 5, 2005**

Paul J. Esatto, Jr.
Reg. No. 30,749
SCULLY, SCOTT, MURPHY & PRESSER
400 Garden City Plaza, Ste. 300
Garden City, NY 11530
(516) 742-4343
PJE:DAT:jam

I hereby certify that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail in an envelope addressed to "Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450" [37 CFR 1.8(a)] on

December 5, 2005
(Date)

Signature of Person Mailing Correspondence

Paul J. Esatto, Jr.

Typed or Printed Name of Person Mailing Correspondence

CC:

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

Applicant:	Yianilos et al.	Examiner:	Etienne Pierre LeRoux
Serial No.:	09/927,589	Art Unit:	2161
Filed:	August 10, 2001	Docket:	13845
For:	A SYNCHRONIZABLE TRANSACTIONAL DATABASE METHOD AND SYSTEM	Dated:	December 5, 2005
Conf. No.:	8179		

Mail Stop Appeal Brief-Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA. 22313-1450

APPEAL BRIEF

CERTIFICATE OF MAILING UNDER 37 C.F.R. §1.8(a)

I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to: Mail Stop Appeal Brief-Patents, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on December 5, 2005.

Dated: December 5, 2005

Paul J. Esatto, Jr.

TABLE OF CONTENTS

I. Real Party in Interest

The real party in interest of the present application is NEC Corporation, the assignee of the entire right, title and interest in the above-identified patent application.

II. Related Appeals and Interferences

No other appeals and interferences are known which directly affect, or will be directly affected by, or have a bearing on, the disposition of the pending appeal.

III. Status of Claims

Claims 1-27 are canceled.

Claim 28 stands rejected under 35 U.S.C. §103(a) based on U.S. Patent No. 6,141,664 issued to Boothby in view of U.S. Patent No. 5,440,732 issued to Lomet et al.

Claim 29 stands rejected under 35 U.S.C. §112, first paragraph and under 35 U.S.C. §103(a) based on U.S. Patent No. 6,141,664 issued to Boothby in view of U.S. Patent No. 5,440,732 issued to Lomet et al.

Claim 30 stands rejected under 35 U.S.C. §103(a) based on U.S. Patent No. 6,141,664 issued to Boothby in view of U.S. Patent No. 5,440,732 issued to Lomet et al.

Claim 31 stands rejected under 35 U.S.C. §103(a) based on U.S. Patent No. 6,141,664 issued to Boothby in view of U.S. Patent No. 5,440,732 issued to Lomet et al. and further in view of U.S. Patent No. 5,668,958 issued to Bendert et al.

Claim 32 stands rejected under 35 U.S.C. §103(a) based on U.S. Patent No. 6,141,664 issued to Boothby in view of U.S. Patent No. 5,440,732 issued to Lomet et al. and further in view of U.S. Patent No. 5,668,958 issued to Bendert et al.

Claim 33 stands rejected under 35 U.S.C. §103(a) based on U.S. Patent No. 6,141,664 issued to Boothby in view of U.S. Patent No. 5,440,732 issued to Lomet et al., further in view of U.S. Patent No. 5,668,958 issued to Bendert et al. and still further in view of U.S. Patent No. 5,089,952 issued to Bozman.

Claim 34 stands rejected under 35 U.S.C. §103(a) based on U.S. Patent No. 6,141,664 issued to Boothby in view of U.S. Patent No. 5,440,732 issued to Lomet et al., further in view of U.S. Patent No. 5,668,958 issued to Bendert et al. and still further in view of U.S. Patent No. 5,778,375 issued to Hecht.

Claim 35 stands rejected under 35 U.S.C. §112, first paragraph and under 35 U.S.C. §103(a) based on U.S. Patent No. 6,141,664 issued to Boothby in view of U.S. Patent No. 5,440,732 issued to Lomet et al.

Claim 36 stands rejected under 35 U.S.C. §112, first paragraph and under 35 U.S.C. §103(a) based on U.S. Patent No. 6,141,664 issued to Boothby in view of U.S. Patent No. 5,440,732 issued to Lomet et al. and further in view of U.S. Patent No. 5,668,958 issued to Bendert et al.

Claim 37 stands rejected under 35 U.S.C. §112, first paragraph and under 35 U.S.C. §103(a) based on U.S. Patent No. 6,141,664 issued to Boothby in view of U.S. Patent No. 5,440,732 issued to Lomet et al.

Claim 38 stands rejected under 35 U.S.C. §112, first paragraph and under 35 U.S.C. §103(a) based on U.S. Patent No. 6,141,664 issued to Boothby in view of U.S. Patent No. 5,440,732 issued to Lomet et al. and further in view of U.S. Patent No. 5,668,958 issued to Bendert et al.

Claim 39 stands rejected under 35 U.S.C. §112, first paragraph and under 35 U.S.C. §103(a) based on U.S. Patent No. 6,141,664 issued to Boothby in view of U.S. Patent No. 5,440,732 issued to Lomet et al. and further in view of U.S. Patent No. 5,668,958 issued to Bendert et al.

Claim 40 stands rejected under 35 U.S.C. §112, first paragraph and under 35 U.S.C. §103(a) based on U.S. Patent No. 6,141,664 issued to Boothby in view of U.S. Patent No. 5,440,732 issued to Lomet et al., further in view of U.S. Patent No. 5,668,958 issued to Bendert et al. and still further in view of U.S. Patent No. 5,089,952 issued to Bozman.

Claim 41 stands rejected under 35 U.S.C. §112, first paragraph and under 35 U.S.C. §103(a) based on U.S. Patent No. 6,141,664 issued to Boothby in view of U.S. Patent No. 5,440,732 issued to Lomet et al., further in view of U.S. Patent No. 5,668,958 issued to Bendert et al. and still further in view of U.S. Patent No. 5,778,375 issued to Hecht.

IV. Status of Amendments

All amendments have been entered and considered by the Examiner prior to filing of the Notice of Appeal by Appellants.

V. Summary of Claimed Subject Matter

The invention with respect to Claim 28 comprises the steps of receiving from a remote transaction database (See: e.g., FIG. 2, ref. 4) a single summary hash computed for database records lying in an interval in the remote transaction database (See: e.g., page 14, lines 6-28); comparing the single summary hash to a local single summary hash computed for database records lying in a same interval in the local transaction database (See: e.g., page 18, lines 3-6); and when the single summary hash does not match the local summary hash, partitioning the interval into at least a first and a second sub-interval and requesting from the remote transaction

database a first sub-interval summary hash and a second sub-interval summary hash (See: e.g., page 18, lines 6-16), the first sub-interval summary hash computed for database records lying in the first sub-interval in the remote transaction database and the second sub-interval summary hash computed for database records lying in the second sub-interval in the remote transaction database, whereby the local transaction database can avoid synchronization of database records in a sub-interval with a local sub-interval summary hash that matches a sub-interval summary hash received from the remote transaction database (See: e.g., page 18, lines 6-7).

The invention with respect to Claim 29 comprises the steps identified with respect to Claim 28 above, and further comprising the step of when the local sub-interval summary hash does not match the sub-interval summary hash received from the remote transaction database, further partitioning the sub-interval into at least a first and a second sub-sub-interval (See: e.g., page 13, lines 27-29) and requesting from the remote transaction database a first and second sub-sub-interval summary hash, the first sub-sub-interval summary hash computed for database records lying in the first sub-sub-interval in the remote transaction database and the second sub-sub-interval summary hash computed for database records lying in the second sub-sub-interval in the remote transaction database (See: e.g., page 13, line 28 – page 14, line 5).

The invention with respect to Claim 30 comprises the steps identified with respect to Claim 28 above, and further comprising the step of when the local sub-interval summary hash does not match the sub-interval summary hash received from the remote transaction database, requesting from the remote transaction database a hash value for each database record lying in the sub-interval (See: e.g., page 13, line 28 – page 14, line 5).

The invention with respect to Claim 31 comprises the steps identified with respect to Claim 28 above, and further, wherein each database record has a hash value and wherein each

summary hash for an interval is computed by combining the hash values for each database record in the interval (See: e.g., page 14, line 9 – page 16, line 24).

The invention with respect to Claim 32 comprises the steps identified with respect to Claim 31 above, and further, wherein the summary hashes for different intervals are stored in a tree structure (See: e.g., page 15, lines 1-11).

The invention with respect to Claim 33 comprises the steps identified with respect to Claim 32 above, and further, wherein the tree structure is a B+ tree structure (See: e.g., page 15, lines 1-2).

The invention with respect to Claim 34 comprises the steps identified with respect to Claim 32 above, and further, wherein the hash values are message digests and wherein the hash values are combined to compute a summary hash by an excusive or (XOR) of the hash values (See: e.g., page 13, lines 5-9).

The invention with respect to Claim 35 comprises a database (See: e.g., FIG. 2, ref. 4); an interval hash value computing module (See: e.g., FIG. 2, ref. 5) coupled to the database (See: e.g., FIG. 2) and configured to compute a summary hash of a plurality of hash values (See: e.g., page 12, lines 23-30), each hash value associated with a database record lying in an interval of the database (See: e.g., page 14, lines 15-19); and a synchronization module (See: e.g., FIG. 2, ref. 3) coupled to the database and to the interval hash value computing module (See: e.g., FIG. 2), the synchronization module configured to identify database records that need synchronization (See: e.g., page 17, lines 6-29) by comparing a summary hash from the interval hash value computing module computed for database records lying in an interval of the database with a remote summary hash received from a remote transaction database (See: e.g., page 18, lines 3-16).

The invention with respect to Claim 36 comprises the elements identified with respect to Claim 35 above, and further, wherein the database (See: e.g., FIG. 2, ref. 4) further comprises a transactional support layer (See: e.g., FIG. 2, ref. 6) configured to support a storage layer (See: e.g., page 5, lines 24-25) of the database, the transactional support layer further comprising shadow blocks which provides for atomized updates to the storage layer (See: e.g., page 5, line 26 – page 6, line 10).

The invention with respect to Claim 37 comprises the elements identified with respect to Claim 35 above, and further, wherein the synchronization module (See: e.g., FIG. 2, ref. 3) is further configured to partition an interval into at least a first and second sub-interval (See: e.g., page 13, lines 27-29) when a summary hash for the interval in the database does not match a remote summary hash so as to seek remote summary hashes for the first and second sub-intervals from the remote transaction database (See: e.g., page 13, line 28 – page 14, line 5).

The invention with respect to Claim 38 comprises the elements identified with respect to Claim 36 above, and further, wherein the interval hash value computing module computes a summary hash for database records lying in an interval in the database by combining the hash values associated with each database record lying in the interval (See: e.g., page 14, line 9 – page 16, line 24).

The invention with respect to Claim 39 comprises the elements identified with respect to Claim 38 above, and further, wherein the summary hashes for different intervals are stored in a tree structure (See: e.g., page 15, lines 1-11).

The invention with respect to Claim 40 comprises the elements identified with respect to Claim 39 above, and further, wherein the tree structure is a B+ tree structure (See: e.g., page 15, lines 1-2).

The invention with respect to Claim 41 comprises the elements identified with respect to Claim 39 above, and further, wherein the hash values are message digests and wherein the hash values are combined to compute a summary hash by using an excusive or (XOR) of the hash values (See: e.g., page 13, lines 5-9).

VI. Grounds of Rejection to be Reviewed on Appeal

1. Rejection under 35 U.S.C. §112, first paragraph.

2. Rejection under 35 U.S.C. §103(a) based on U.S. Patent No. 6,141,664 issued to Boothby in view of U.S. Patent No. 5,440,732 issued to Lomet et al.

3. Rejection under 35 U.S.C. §103(a) based on U.S. Patent No. 6,141,664 issued to Boothby in view of U.S. Patent No. 5,440,732 issued to Lomet et al. and further in view of U.S. Patent No. 5,668,958 issued to Bendert et al.

4. Rejection under 35 U.S.C. §103(a) based on U.S. Patent No. 6,141,664 issued to Boothby in view of U.S. Patent No. 5,440,732 issued to Lomet et al., further in view of U.S. Patent No. 5,668,958 issued to Bendert et al. and still further in view of U.S. Patent No. 5,089,952 issued to Bozman.

5. Rejection under 35 U.S.C. §103(a) based on U.S. Patent No. 6,141,664 issued to Boothby in view of U.S. Patent No. 5,440,732 issued to Lomet et al., further in view of U.S. Patent No. 5,668,958 issued to Bendert et al. and still further in view of U.S. Patent No. 5,778,375 issued to Hecht.

VII. Argument

1. Rejection Under 35 U.S.C. §112, First Paragraph

A. Claim 29

The terms "a first" and "a second sub-sub-interval," and "a first" and "a second sub-sub-interval summary hash" are alleged to be unsupported. As explained in Appellants' disclosure, page 14:

> This synchronization operation works in a number of communication rounds. In each round, the key range of interest is partitioned into smaller sub-ranges. For each sub-range, the two databases compute the summary of records lying in that sub-range and one of the databases sends its summaries to the other side. The corresponding summaries from the two sides are compared and the operation is recursively applied to sub-ranges whose summaries do not match. Only those records are transferred from one side to the other which (1) are missing on the other side, or (2) have a mismatching record on the other side. Thus, unnecessary transfer of large amounts of data is prevented.

Thus, the synchronization operation is performed by comparing or by matching a hash value obtained for an interval of data of a remote database with a hash value for a corresponding interval of a second database. If the hash values match, then the data contained in the interval and the corresponding interval do not need to be synchronized. However, if the hashes do not match, then a hash of a sub-interval of the remote database is compared with a hash of a corresponding sub-interval of the second database. This is what is meant by the operation being recursively applied to sub-ranges if summaries or hashes do not match. If these match, then a second sub-

interval of the same interval of the remote database is compared with a hash of a corresponding second sub-interval of the second database. However, if the hash of the sub-interval and the corresponding sub-interval do not match, then a hash of the sub-sub interval and a corresponding sub-sub-interval are compared. This recursive process is repeated until the smallest range (which is separately accessible, for example, a data block representing a set of records, a record, or a field of a record) for which the hashes do not match is reached (or until some other desired size range of data is reached).

It is respectfully submitted that a person of ordinary skill in the art would have readily understood each of the foregoing terms based on Appellants' disclosure, and therefore this rejection should now be withdrawn.

B. Claims 35-41

(1) The term "interval hash value computing module" is alleged to be unsupported. However, Appellants' FIG. 2 provides a summarizable database containing a BxTree 5 that "implements a B+-tree based database engine augmented with functions for supporting range synchronization. For example, a typical operation on the database would input a key range (i.e. an interval of the space of keys) and return a short summary. The summary would typically contain a digest of the records of interest of all the records in the database whose keys lie in the given key range." (See: page 12, lines 23-27). The BxTree 5 performs a "Get Interval Hash" function as described on page 14, lines 21-28 of the disclosure as follows:

> **Get_Interval_Hashes**: The input is an interval *I* of *K* and a positive integer *H*. The function partitions *I* into at most *H* disjoint sub-intervals and returns a list of triplets of the form (key_interval, num_records, hash). The list has one triplet for each sub-interval. The first element of the triplet is the sub-interval itself; the second and third elements are, respectively, the number and a fixed size digest of all the records in the database whose key fields belong to the sub-interval.

> Whether the database has any records with key field belonging to *I* or not, the list returned is always non-empty and the sub-intervals in the list form a disjoint partition of *I*. (Italics in the original.) Thus, the interval hash value computing module (or BxTree) computes a hash or value that serves as a "summary" of the interval, i.e., the range of data contained in the database.

(2) The term "a synchronization module coupled to the database and to the interval hash value computing module" is alleged to be unsupported. However, Appellants' FIG. 2 clearly identifies a Synchronization Facility 3 (OSynch). The OSynch module, according to an aspect of Appellants' claimed invention, implements range synchronization for the BxTree database embodiment of Appellants' claimed invention. The interval hash value computing module (i.e. BxTree module) and the OSynch module "talk to each other through a very clean interface..." (See: page 12, lines 20-22). Therefore, in order for the OSynch module to talk to the BxTree module, the two modules would inherently need to be coupled. Additionally, as the OSynch module is for synchronizing two databases, the module would also need to be coupled to the databases to be synchronized. For example, pages 13-14 provide as follows:

> The osynch (or object synchronization) module, implements the range synchronization operation on bxtree databases... For each sub-range, the two databases compute the summary of records lying in that sub-range and one of the databases sends its summaries to the other side. The corresponding summaries from the two sides are compared and the operation is recursively applied to sub-ranges whose summaries do not match. Only those records are transferred from one side to the other which (1) are missing on the other side, or (2) have a mismatching record on the other side. Thus, unnecessary transfer of large amounts of data is prevented.

2. Rejection under 35 U.S.C. §103(a)

A. Claim 28-30, 35 and 37

Claims 28-30, 35 and 37 are rejected under 35 U.S.C. § 103 as allegedly being obvious over Boothby, U.S. Patent No. 6,141,664 in view of Lomet et al., U.S. Patent No. 5,440,732.

Among the problems recognized and solved by Appellants' claimed invention is that of having to compare most of the data of two databases to synchronize the databases, or having to rewrite entire files specifically provided, in order to synchronize the databases. According to an aspect of Appellants' claimed invention, hash values that "summarize" an interval of the first database and the second database are obtained and these hash values are compared or matched to determine whether a synchronization of that interval is necessary. If the hash values for the interval do not match, then as discussed above with respect to the synchronization facility, sub-interval hash values may be obtained and the process repeated recursively.

For at least the following reasons, Appellants' claimed invention is neither anticipated by nor obvious from the cited references. By way of example, independent claim 28 requires synchronizing based on whether the single summary hash for the records of the interval match the local summary hash of the interval. Further, independent claim 35 requires identifying database records that need synchronizing by comparing the summary hash for the interval with a remote summary hash for the interval.

Boothby is directed to synchronization of incompatible databases based on a date range (Boothby, Abstract); such that incompatible databases (databases with different types of data structures, Boothby, column 1, lines 39-50) are synchronized without overloading the storage capacity of a database contained in a smaller device (Boothby, column 2, lines 1-10). Boothby discloses that the synchronization of the two databases is performed for the current date range and a prior date range (Boothby, column 2, lines 13-23); such that the user specifies the date

range for which synchronization is desired when an "incremental" date range synchronization is performed (Boothby, column 5, lines 42-55); or a synchronization from scratch for the entire date range may be performed (Boothby, column 5, lines 5-20).

Boothby does not disclose or suggest obtaining a summary hash value for a record or set of data for an interval of a database, as *inter alia*, required by independent claims 28 and 35. Nor does Boothby disclose or suggest synchronizing databases based on a comparison or matching of such summary hash values, as further required by independent claims 28 and 35. Additionally, the Boothby-disclosed incremental synchronization does not perform a comparison between two databases in an initial interval and then, if the two databases do not match, partitioning the interval into a first sub-interval and a second sub-interval for performing further comparisons of the two databases in narrower intervals within the initial interval.

Lomet discloses a locking strategy for a database employing an index tree (Lomet, Abstract; column 1, lines 6-8) for effective locking and unlocking of a database resource while an update is in process (Lomet, column 3, line 35 - column 4, line 47). Lomet discloses a database indexing approach using a B-tree organization indexing approach, and Lomet discloses a search mode using a hash function, in which a hash value calculated for a search key is used to compute the address of the "bucket" in which the record resides that contains the searched for key value (Lomet, column 40, lines 40-63). Thus, Lomet describes that a hashing algorithm applied to the search key value yields an address of the bucket in which the desired record resides.

Lomet does not disclose or suggest synchronizing database records based on whether the summary hash for the record of the interval matches the local summary hash of the interval, as *inter alia* required by independent claim 28. Further, Lomet does not disclose or suggest identifying database records that need synchronizing by comparing the summary hash for the

interval with a remote summary hash for the interval, as *inter alia* required by independent claim 35. The hash value disclosed in Lomet is for random access of a single file and not a summary hash as recited in the claims. In fact, Lomet specifically states that use of hash-type indexes for range searching is unpopular because their performance in range searching is abysmal. (See: Lomet col. 9, lines 40-63). Consequently, Lomet teaches away from using a summary hash for performing range searching as recited in Appellant's claims. Therefore, one skilled in the art would not look to Lomet for synchronizing database records based on whether the summary hash for the record of the interval matches the local summary hash of the interval.

First, Lomet does not disclose or suggest obtaining a summary hash value for a set of data of an interval of a database. As discussed, the hashing function is applied to a search key and yields an address for a record. However, Lomet does not disclose or suggest a hash value obtained for a set of data or data record(s).

Moreover, since Lomet does not disclose or suggest this feature, Lomet is incapable of disclosing or suggesting synchronizing databases based on a comparison or matching of such summary hash values, as *inter alia*, required by independent claims 28 and 35. Therefore, Boothby and Lomet, even if combined and taken together as a whole, do not disclose or suggest the recitation of independent claims 28 and 35.

More generally, Boothby and Lomet belong to the conventional art recognized by Appellants' claimed invention, because Boothby and Lomet do not disclose or suggest the above-discussed problems recognized and solved by Appellants' claimed invention. For instance, the problem of having to compare most of the data of two databases to synchronize the databases, or having to rewrite entire files specified in advance in order to synchronize the databases, is not disclosed or suggested by Boothby and Lomet. Also, at least one of the solutions provided by

Appellants' claimed invention – synchronizing database data based on hash value summaries for data contained in corresponding intervals of the databases – is not disclosed. Accordingly, the hypothetical combination of Boothby and Lomet would not even remotely disclose or suggest Appellants' claimed invention as recited in Claims 28-30, 35 and 37.

B. Claims 31, 32, 36, 38 and 39

Claims 31, 32, 36, 38 and 39 are rejected under 35 U.S.C. § 103 as allegedly being obvious over Boothby and Lomet in view of Bendert et al., U.S. Patent No. 5,668,958.

Bendert does not cure the above-discussed deficiencies of the combination of Boothby and Lomet. Therefore, since claims 31 and 32 depend from independent claim 28, and claims 36, 38 and 39 depend from independent claim 35, claims 31, 32, 36, 38 and 39 are patentably distinguishable over the prior art for at least the reasons that independent claims 28 and 35 are patentably distinguishable over the prior art.

C. Claims 33 and 40

Claims 33 and 40 are rejected under 35 U.S.C. § 103 as being obvious from Boothby, Lomet and Bendert in view Bozman, U.S. Patent No. 5,089,952.

Bozman does not cure the above-discussed deficiencies of the combination of Boothby and Lomet. Therefore, since claims 33 and 40 depend from independent claims 28 and 35, respectively, claims 33 and 40 are patentably distinguishable over the prior art for at least the reasons that independent claims 28 and 35 are patentably distinguishable over the prior art.

D. Claims 34 and 41

Claims 34 and 41 are rejected under 35 U.S.C. § 103 as allegedly being obvious from Boothby, Lomet and Bendert in view of Hecht, U.S. Patent No. 5,778,375. This rejection is traversed.

Hecht does not cure the above-discussed deficiencies of the combination of Boothby and Lomet. Therefore, since claims 34 and 41 depend from independent claims 28 and 35, respectively, claims 34 and 41 are patentably distinguishable over the prior art for at least the reasons that independent claims 28 and 35 are patentably distinguishable over the prior art. Therefore, this rejection should now be withdrawn.

3. Conclusion

It is clear that all of the limitations of claims 28-41 are not taught or suggested by the references of Boothby, Lomet, Bendert, Bozman and Hecht, individually or in any proper combination. Therefore, Appellant respectfully submits that the Examiner has not met his burden of establishing a prima facie case of obviousness based on the prior art, as required by 35 U.S.C. §103(a)[1]. No objective teaching in Boothby, Lomet, Bendert, Bozman and Hecht, individually or in any proper combination, would lead an individual of ordinary skill in the art to make the present invention. Accordingly, Appellant respectfully requests that the rejections made in the Final Rejection dated May 5, 2005, and in the Advisory Action of September 13, 2005, be reversed by the Board of Patent Appeals and Interferences.

Respectfully submitted,

Paul J. Esatto, Jr.
Registration No. 30,749

SCULLY, SCOTT, MURPHY & PRESSER
400 Garden City Plaza - Ste. 300
Garden City, New York 11530
(516) 742-4343

PJE:DAT:jam

1 In re Piasecki, 745 F.2d 1468, 1471-72, 223 USPQ 785, 787-88 (Fed.Cir. 1984).

VIII. CLAIMS APPENDIX

28. A method of synchronizing a local transaction database with a remote transaction database, comprising:

receiving from the remote transaction database a single summary hash computed for database records lying in an interval in the remote transaction database;

comparing the single summary hash to a local single summary hash computed for database records lying in a same interval in the local transaction database; and

when the single summary hash does not match the local summary hash, partitioning the interval into at least a first and a second sub-interval and requesting from the remote transaction database a first sub-interval summary hash and a second sub-interval summary hash, the first sub-interval summary hash computed for database records lying in the first sub-interval in the remote transaction database and the second sub-interval summary hash computed for database records lying in the second sub-interval in the remote transaction database,

whereby the local transaction database can avoid synchronization of database records in a sub-interval with a local sub-interval summary hash that matches a sub-interval summary hash received from the remote transaction database.

29. The method of claim 28 further comprising the step of: when the local sub-interval summary hash does not match the sub-interval summary hash received from the remote transaction database, further partitioning the sub-interval into at least a first and a second sub-sub-interval and requesting from the remote transaction database a first and second sub-sub-interval summary hash, the first sub-sub-interval summary hash computed for database records lying in the first sub-sub-interval in the remote transaction database and the second sub-sub-interval summary hash computed for database records lying in the second sub-sub-interval in the remote transaction database.

30. The method of claim 28 further comprising the step of: when the local sub-interval summary hash does not match the sub-interval summary hash received from the remote transaction database, requesting from the remote transaction database a hash value for each database record lying in the sub-interval.

31. The method of claim 28, wherein each database record has a hash value and wherein each summary hash for an interval is computed by combining the hash values for each database record in the interval.

32. The method of claim 31, wherein the summary hashes for different intervals are stored in a tree structure.

33. The method of claim 32, wherein the tree structure is a B+ tree structure.

34. The method of claim 32, wherein the hash values are message digests and wherein the hash values are combined to compute a summary hash by an excusive or (XOR) of the hash values.

35. A synchronizable transactional database comprising:

a database;

an interval hash value computing module coupled to the database and configured to compute a summary hash of a plurality of hash values, each hash value associated with a database record lying in an interval of the database; and

a synchronization module coupled to the database and to the interval hash value computing module, the synchronization module configured to identify database records that need synchronization by comparing a summary hash from the interval hash value computing module computed for database records lying in an interval of the database with a remote summary hash received from a remote transaction database.

36. The system of claim 35, wherein the database further comprises a transactional support layer configured to support a storage layer of the database, the transactional support layer further comprising shadow blocks which provides for atomized updates to the storage layer.

37. The system of claim *35,* wherein the synchronization module is further configured to partition an interval into at least a first and second sub-interval when a summary hash for the interval in the database does not match a remote summary hash so as to seek remote summary hashes for the first and second sub-intervals from the remote transaction database.

38. The system of claim 36, wherein the interval hash value computing module computes a summary hash for database records lying in an interval in the database by combining the hash values associated with each database record lying in the interval.

39. The system of claim 38, wherein the summary hashes for different intervals are stored in a tree structure.

40. The system of claim 39, wherein the tree structure is a B+ tree structure.

41. The system of claim 39, wherein the hash values are message digests and wherein the hash values are combined to compute a summary hash by using an excusive or (XOR) of the hash values.

IX. EVIDENCE APPENDIX

None

X. RELATED PROCEEDINGS APPENDIX

None